arXiv:1702.07811v1 [cs.LG] 25 Feb 2017

Adaptive Neural Networks for Fast Test-Time Prediction

Tolga Bolukbasi [1] **Joseph Wang** [2] **Ofer Dekel** [3] **Venkatesh Saligrama** [1]

Abstract

We present an approach to adaptively utilize deep neural networks in order to reduce the evaluation time on new examples without loss of classification performance. Rather than attempting to redesign or approximate existing networks, we propose two schemes that adaptively utilize networks. First, we pose an adaptive network evaluation scheme, where we learn a system to adaptively choose the components of a deep network to be evaluated for each example. By allowing examples correctly classified using early layers of the system to exit, we avoid the computational time associated with full evaluation of the network. Building upon this approach, we then learn a network selection system that adaptively selects the network to be evaluated for each example. We exploit the fact that many examples can be correctly classified using relatively efficient networks and that complex, computationally costly networks are only necessary for a small fraction of examples. By avoiding evaluation of these complex networks for a large fraction of examples, computational time can be dramatically reduced. Empirically, these approaches yield dramatic reductions in computational cost, with up to a 2.8x speedup on state-of-the-art networks from the ImageNet image recognition challenge with minimal (less than 1%) loss of accuracy.

1. Introduction

Deep neural networks (DNNs) are among the most powerful and versatile machine learning techniques, achieving state-of-the-art accuracy in a variety of important applications, such as visual object recognition (He et al., 2016), speech recognition (Hinton et al., 2012), and machine translation (Sutskever et al., 2014). However, the power of DNNs comes at a considerable cost, namely, the computational cost of applying them to new examples. This cost, often called the *test-time cost*, has increased rapidly for many tasks (see Fig. 1) with ever-growing demands for improved performance in state-of-the-art systems. As a point of fact, the *Resnet152* (He et al., 2016) architecture with 152 layers, realizes a substantial 4.4% accuracy gain in top-5 performance over GoogLeNet (Szegedy et al., 2015) on the large-scale ImageNet dataset (Russakovsky et al., 2015) but is about 14X slower at test-time. The high test-time cost of state-of-the-art DNNs means that they can only be deployed on powerful computers, equipped with massive GPU accelerators. As a result, technology companies spend billions of dollars a year on expensive and power-hungry computer hardware. Moreover, high test-time cost prevents DNNs from being deployed on resource constrained platforms, such as those found in Internet of Things (IoT) devices, smart phones, and wearables. This problem has given rise to a concentrated research effort to reduce the test-time cost of DNNs. Most of the work on this topic focuses on designing more efficient network topologies and on compressing pre-trained models using various techniques (see related work below). We propose a different approach, which leaves the original DNN intact and instead changes the way in which we apply the DNN to new examples. We exploit the fact that natural data is typically a

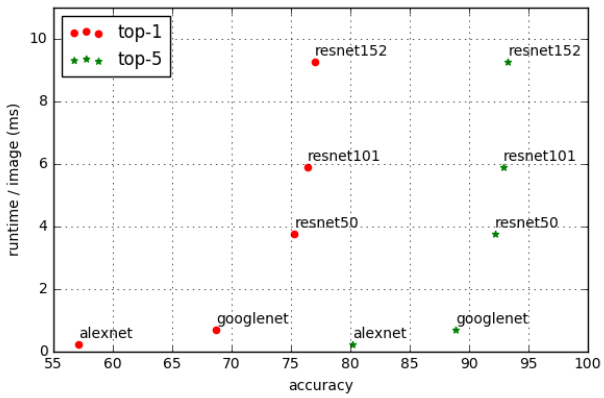

Figure 1. Performance versus evaluation complexity of the DNN architectures that won the ImageNet challenge over past several years. The model evaluation times increase exponentially with respect to the increase in accuracy.

[1]Boston University, Boston, MA, USA [2]Amazon, Cambridge, MA, USA [3]Microsoft Research, Redmond, WA, USA. Correspondence to: Tolga Bolukbasi <tolgab@bu.edu>.

mix of easy examples and difficult examples, and we posit that the easy examples do not require the full power and complexity of a massive DNN.

We pursue two concrete variants of this idea. First, we propose an adaptive early-exit strategy that allows easy examples to bypass some of the network's layers. Before each expensive neural network layer (e.g., convolutional layers), we train a policy that determines whether the current example should proceed to the next layer, or be diverted to a simple classifier for immediate classification. Our second approach, an adaptive network selection method, takes a set of pre-trained DNNs, each with a different cost/accuracy trade-off, and arranges them in a cascade (Trapeznikov & Saligrama, 2013), with the the cheapest model first and the most expensive one last. We then train an exit policy before each model in the cascade, which determines whether we should rely on the current model's predictions or forward the example to the next model in the cascade.

We demonstrate the merits of our techniques on the ImageNet object recognition task, using a number of popular pretrained DNNs. The early exit technique speeds up the average test-time evaluation of GoogLeNet (Szegedy et al., 2015), and Resnet50 (He et al., 2016) by 20-30% within reasonable accuracy margins. The network cascade achieves 2.8x speed-up compared to pure Resnet50 model at 1% top-5 accuracy loss and 1.9x speed-up with no change in model accuracy. We also show that our method can approximate a Oracle-policy that can see true errors suffered for each instance.

In addition to reducing the average test-time cost of DNNs, it is worth noting that our techniques are compatible with the common design of large systems of mobile devices, such as smart phone networks or smart surveillance-camera networks. These systems typically include a large number of resource-constrained edge devices that are connected to a central and resource-rich cloud. One of the main challenges involved in designing these systems is determining whether the machine-learned models will run in the devices or in the cloud. Offloading all of the work to the cloud can be problematic due to network latency, limited cloud ingress bandwidth, cloud availability and reliability issues, and privacy concerns. Our approach can be used to design such a system, by deploying a small inaccurate model and an exit policy on each device and a large accurate model in the cloud. Easy examples would be handled by the devices, while difficult ones would be forwarded to the cloud. Our approach naturally generalizes to a fog computing topology (where resource constrained edge devices are connected to a more powerful local gateway computer, which in turn is connected to a sequence of increasingly powerful computers along the path to the data-center).

2. Related Work

Past work on reducing evaluation time of deep neural networks has centered on reductions in precision and arithmetic computational cost, design of efficient network structure, and compression or sparsification of networks to reduce the number of convolutions, neurons, and edges. The approach proposed in this paper is complimentary. Our approach does not modify network structure or training and can be applied in tandem with these approaches to further reduce computational cost.

The early efforts to compress large DNNs used a large *teacher model* to generate an endless stream of labeled examples for a smaller *student model* (BuciluǍ et al., 2006; Hinton et al., 2015). The wealth of labeled training data generated by the teacher model allowed the small student model to mimic its accuracy.

Reduced precision networks (Gong et al., 2014; Courbariaux et al., 2015; Chen et al., 2015; Hubara et al., 2016a; Wu et al., 2016; Rastegari et al., 2016; Hubara et al., 2016b) have been extensively studied to reduce the memory footprint of networks and their test-time cost. Similarly, computationally efficient network structures have also been proposed to reduce the computational cost of deep networks by exploiting efficient operations to approximate complex functions, such as the inception layers introduced in GoogLeNet (Szegedy et al., 2015).

Network sparsification teachniques attempt to identify and prune away redundant parts of a large neural networks. A common approach is to remove unnecessary nodes/edges from the network(Liu et al., 2015; Iandola et al., 2016; Wen et al., 2016). In convolutional neural networks, the expensive convolution layers can be approximated (Bagherinezhad et al., 2016) and redundant computation can be avoided (Figurnov et al., 2016).

Our techniques for adaptively computing DNNs borrow ideas from the related sensor selection problem (Xu et al., 2013; Kusner et al., 2014; Wang et al., 2014; 2015). The goal of sensor selection is to adaptively choose sensor measurements or features for each example.

3. Adaptive Early Exit Networks

Our first approach to reducing the test-time cost of deep neural networks is an early exit strategy. We denote a labeled example as $(x, y) \in \mathbb{R}^d \times \{1, \ldots, \mathcal{L}\}$, where d is the dimension of the data and $\{1, \ldots, \mathcal{L}\}$ is the set of classes represented in the data. We define the distribution generating the examples as $\mathcal{X} \times \mathcal{Y}$. For a predicted label $\hat{y}$, we denote the loss $L(\hat{y}, y)$. In this paper, we focus on the task of classification and therefore focus on the indicator loss $L(\hat{y}, y) = \mathbb{1}_{\hat{y}=y}$, though the proposed approach can be ap-

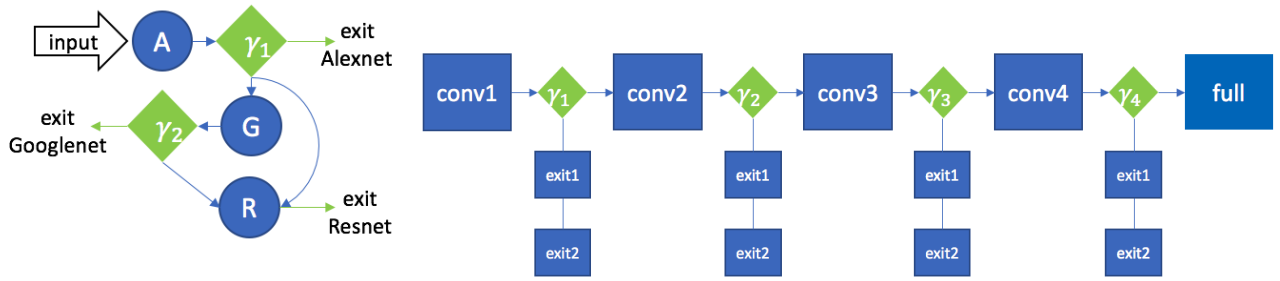

Figure 2. **(Left)** An example network selection system topology for networks Alexnet(A), GoogLeNet(G) and Resnet(R). Green γ blocks denote the selection policy. The policy evaluates Alexnet, receives confidence feedback and decides to jump directly to Resnet or send the sample to GoogLeNet->Resnet cascade. **(Right)** An example early exit system topology (based on Alexnet). The policy chooses one of the multiple exits available to it at each stage for feedback. If the sample is easy enough, the system sends it down to exit, otherwise it sends the sample to the next layer.

plied using any loss function, such as the ℓ_2-norm for regression.

As a running DNN example, we consider the AlexNet architecture (Krizhevsky et al., 2012), which is composed of 5 convolutional layers followed 3 fully connected layers. During evaluation of the network, computing each convolutional layer takes more than 3 times longer than computing a fully connected layer, so we consider a system that allows an example to exit the network after each of the first 4 convolutional layers. Let $\hat{y}(x)$ denote the label predicted by the network for example x and assume that computing this prediction takes a constant time of T. Moreover, let $\sigma_k(x)$ denote the output of the k^{th} convolutional layer for example x and let t_k denote the time it takes to compute this value (from the time that x is fed to the input layer). Finally, let $\hat{y}_k(x)$ be the predicted label if we exit after the k^{th} layer.

After computing the k^{th} convolutional layer, we introduce a decision function γ_k that determines whether the example should exit the network with a label of $\hat{y}_k(x)$ or proceed to the next layer for further evaluation. The input to this decision function is the output of the corresponding convolutional layer $\sigma_k(x)$, and the value of $\gamma_k(\sigma_k(x))$ is either -1 (indicating an early exit) or 1. This architecture is depicted on the right-hand side of Fig. 2.

For concreteness, we focus on the problem of learning the decision function γ_4, which determines if an example should exit after the fourth convolutional layer or whether it will be classified using the entire network. The time it takes to predict the label of example x depends on this decision and can be written as

$$T_4\left(x,\gamma_4\right)=\begin{cases}T+\tau(\gamma_4) & \text{if } \gamma_4(\sigma_4(x))=1\\ t_4+\tau(\gamma_4) & \text{otherwise}\end{cases}, \quad (1)$$

where $\tau(\gamma_4)$ is the computational time required to evaluate the function γ_4. Our goal is to learn a system that trades-off the evaluation time and the induced error:

$$\underset{\gamma_4\in\Gamma}{\operatorname{argmin}}\,\mathbb{E}_{x\sim\mathcal{X}}\left[T_4(x,\gamma_4)\right]+\lambda\mathbb{E}_{(x,y)\sim\mathcal{X}\times\mathcal{Y}}\Big[\Big(L\left(\hat{y}_4(x),y\right) - L\left(\hat{y}(x),x\right)\Big)_+\mathbb{1}_{\gamma_4(\sigma_4(x))=-1}\Big] \quad (2)$$

where $(\cdot)_+$ is the function$(z)_+=\max(z,0)$ and $\lambda\in\mathbb{R}^+$ is a trade-off parameter that balances between evaluation time and error. Note that the function $T_4\left(x,\gamma_4\right)$ can be expressed as a sum of indicator functions:

$$\begin{aligned}T_4\left(x,\gamma_4\right)=&\left(T+\tau(\gamma_4)\right)\mathbb{1}_{\gamma_4(\sigma_4(x))=1}\\&+\left(t_4+\tau(\gamma_4)\right)\mathbb{1}_{\gamma_4(\sigma_4(x))=-1}\\=&T\mathbb{1}_{\gamma_4(\sigma_4(x))=1}+t_4\mathbb{1}_{\gamma_4(\sigma_4(x))=-1}+\tau_4(\gamma_4)\end{aligned}$$

Substituting for $T_4(x,\gamma_4)$ allows us to reduce the problem to an importance weighted binary learning problem:

$$\underset{\gamma_4\in\Gamma}{\operatorname{argmin}}\,\mathbb{E}_{(x,y)\sim\mathcal{X}\times\mathcal{Y}}\left[C_4(x,y)\mathbb{1}_{\gamma_4(\sigma_4(x))\neq\beta_4(x)}\right]+\tau(\gamma_4) \quad (3)$$

where $\beta_4(x)$ and $C_4(x,y)$ are the optimal decision and cost at stage 4 for the example (x,y) defined:

$$\beta_4(x)=\begin{cases}-1 & \text{if } T>\Big(t_4+\lambda\Big(L\left(\hat{y}_4(x),y\right)\\ & \qquad -L\left(\hat{y}(x),x\right)\Big)_+\Big)\\ 1 & \text{otherwise}\end{cases}$$

and

$$C_4(x,y)=\left|T-t_4-\lambda\left(L\left(\hat{y}_4(x),y\right)-L\left(\hat{y}(x),x\right)\right)_+\right|.$$

Note that the regularization term, $\tau(\gamma_4)$, is important to choose the optimal functional form for the function γ_4

as well as a natural mechanism to define the structure of the early exit system. Rather than limiting the family of function Γ to a single functional form such as a linear function or a specific network architecture, we assume the family of functions Γ is the union of multiple functional families, notably including the constant decision functions $\gamma_4(x) = 1, \forall x \in |\mathcal{X}|$. Although this constant function does not allow for adaptive network evaluation at the specific location, it additionally does not introduce any computational overhead, that is, $\tau(\gamma_4) = 0$. By including this constant function in Γ, we guarantee that our technique can only decrease the test-time cost.

Empirically, we find that the most effective policies operate on classifier confidences such as classification entropy. Specifically, we consider the family of functions Γ as the union of three functional families, the aforementioned constant functions, confidence features generated from linear classifiers applied to $\sigma_4(x)$, and confidence features generated from deep classifiers applied to $\sigma_4(x)$.

Rather than optimizing jointly over all three networks, we leverage the fact that the optimal solution to Eqn. (3) can be found by optimizing over each of the three families of functions independently. For each family of functions, the policy evaluation time $\tau(\gamma_4)$ is constant, and therefore solving (3) over a single family of functions is equivalent to solving an unregularized learning problem. We exploit this by solving the three unregularized learning problems and taking the minimum over the three solutions.

In order to learn the sequence of decision functions, we consider a bottom-up training scheme, as previously proposed in sensor selection (Wang et al., 2015). In this scheme, we learn the deepest early exit function first, then fix the outputs. Fixing the outputs of this trained function, we then train the early exit function immediately preceding the deepest early exit function (γ_3 in Fig. 2).

For a general early exit system, we recursively define the future time, $T_k(x, \gamma_k)$, and the future predicted label, $\tilde{y}_k(x, \gamma_k)$, as

$$T_k(x, \gamma_k) = \begin{cases} T + \tau(\gamma_k) & \text{if } \gamma_k(\sigma_k(x)) = 1 \text{ and } k = K \\ T_{k+1}(x, \gamma_k + 1) & \text{if } \gamma_k(\sigma_k(x)) = 1 \text{ and } k < K \\ \quad +\tau(\gamma_k) & \\ t_k + \tau(\gamma_k) & \text{otherwise} \end{cases}$$

and

$$\tilde{y}_k(x, \gamma_k) = \begin{cases} \hat{y}(x) & \text{if } k = K + 1 \\ \hat{y}(x) & \text{if } k = K \\ & \quad \text{and } \gamma_k(\sigma_k(x)) = 1 \\ \tilde{y}_{k+1}(x, \gamma_{k+1}) & \text{if } k < K \\ & \quad \text{and } \gamma_k(\sigma_k(x)) = -1 \\ \hat{y}_k(x) & \text{otherwise} \end{cases}.$$

Using these definitions, we can generalize Eqn. (3). For a system with K early exit functions, the k^{th} early exit function can be trained by solving the supervised learning problem:

$$\underset{\gamma_k \in \Gamma}{\operatorname{argmin}} \mathbb{E}_{(x,y)\sim\mathcal{X}\times\mathcal{Y}} \left[C_k(x, y) \mathbb{1}_{\gamma_k(x) \neq \beta_k(\sigma_k(x))} \right] + \tau(\gamma_k), \tag{4}$$

where optimal decision and cost $\beta_k(x)$ and $C_k(x, y)$ can be defined:

$$\beta_k(x) = \begin{cases} -1 & \text{if } k < K \text{ and } T_{k+1}(x, \gamma_{k+1}) \geq t_k + \\ & \quad \lambda \left(L\left(\hat{y}_k(x), y\right) - L\left(\tilde{y}_{k+1}(x), x\right) \right)_+ \\ -1 & \text{if } k = K \text{ and } T \geq t_k + \\ & \quad \lambda \left(L\left(\hat{y}_k(x), y\right) - L\left(\tilde{y}_{k+1}(x), x\right) \right)_+ \\ 1 & \text{otherwise} \end{cases}$$

and

$$C_k(x, y) = \begin{cases} \Big| T_{k+1}(x, \gamma_{k+1}) - t_k & \text{if } k < K \\ -\lambda \Big(L\left(\hat{y}_k(x), y\right) - L\left(\tilde{y}_{k+1}(x), x\right) \Big)_+ \Big| & \\ \Big| T - t_k & \text{otherwise} \\ -\lambda \left(L\left(\hat{y}_k(x), y\right) - L\left(\hat{y}(x), x\right) \right)_+ \Big| & \end{cases}.$$

Eqn. (4) allows for efficient training of an early exit system by sequentially training early exit decision functions from the bottom of the network upwards. Furthermore, by including constant functions in the family of functions Γ and training early exit functions in all potential stages of the system, the early exit architecture can also naturally be discovered.

4. Network Selection

As shown in Fig. 1, the computational time has grown dramatically with respect to classification performance. Rather than attempting to reduce the complexity of the state-of-the-art networks, we instead leverage this nonlinear growth by extending the early exiting strategy to the regime of network selection. Conceptually, we seek to exploit the fact that many examples are correctly classified by relatively efficient networks such as alexnet and googlenet, whereas only a small fraction of examples are correctly classified by computationally expensive networks such as resnet 152 and incorrectly classified by googlenet and alexnet.

As an example, assume we have three pre-trained networks, N_1, N_2, and N_3. For an example x, denote the predictions for the networks as $N_1(x)$, $N_2(x)$, and $N_3(x)$. Additionally, denote the evaluation times for each of the networks as $\tau(N_1)$, $\tau(N_2)$, and $\tau(N_3)$.

As in Fig. 2, the adaptive system composed of two decision functions that determine which networks are evaluated for each example. First, $\kappa_1 : |\mathcal{X}| \rightarrow \{N_1, N_2, N_3\}$ is applied after evaluation of N_1 to determine if the classification decision from N_1 should be returned or if network N_2 or network N_3 should be evaluated for the example. Foe examples that are evaluated on N_2, $\kappa_2 : |\mathcal{X}| \rightarrow \{N_2, N_3\}$ determines if the classification decision from N_2 should be returned or if network N_3 should be evaluated for the example.

Our goal is to learn the functions κ_1 and κ_2 that minimize the average evaluation time subject to a constraint on the average loss induced by adaptive network selection. As in the adaptive early exit case, we first learn κ_2 to trade-off between the average evaluation time and induced error:

$$\min_{\kappa_2 \in \Gamma} \mathbb{E}_{x\sim\mathcal{X}} \left[\tau(N_3)\mathbb{1}_{\kappa_2(x)=N_3}\right] + \tau(\kappa_2) + \lambda \mathbb{E}_{(x,y)\sim\mathcal{X}\times\mathcal{Y}} \left[\Big(L\left(N_2(x), y\right) - L\left(N_3(x), x\right) \Big)_+ \mathbb{1}_{\kappa_2(x)=N_2} \right], \quad (5)$$

where $\lambda \in \mathbb{R}^+$ is a trade-off parameter. As in the adaptive network usage case, this problem can be posed as an importance weighted supervised learning problem:

$$\min_{\kappa_2 \in \Gamma} \mathbb{E}_{(x,y)\sim\mathcal{X}\times\mathcal{Y}} \left[W_2(x,y)\mathbb{1}_{\kappa_2(x)\neq\theta_2(x)}\right] + \tau(\kappa_2), \quad (6)$$

where $\theta_2(x)$ and $W_2(x, y)$ are the cost and optimal decision at stage 4 for the example/label pair (x, y) defined:

$$\theta_2(x) = \begin{cases} N_2 & \text{if } \tau(N_3) > \lambda \left(L\left(N_3(x), y\right) - L\left(N_2(x), x\right)\right)_+ \\ N_3 & \text{otherwise} \end{cases}$$

and

$$W_2(x,y) = \left| \tau(N_3) - \lambda \left(L\left(N_2(x), y\right) - L\left(N_3(x), x\right)\right)_+ \right|.$$

Once κ_2 has been trained according to Eqn. (6), the training times for examples that pass through N_2 and are routed by κ_2 can be defined $T_{\kappa_2}(x) = \tau(N_2) + \tau(\kappa_2) + \tau(N_3)\mathbb{1}_{\kappa_2(x)=N_3}$. As in the adaptive early exit case, we train and fix the last decision function, κ_2, then train the earlier function, κ_1. As before, we seek to trade-off between evaluation time and error:

$$\min_{\kappa_1 \in \Gamma} \mathbb{E}_{x\sim\mathcal{X}} \left[\tau(N_3)\mathbb{1}_{\kappa_1(x)=N_3} + \tau(N_2)\mathbb{1}_{\kappa_1=N_2}\right] + \tau(\kappa_1) + \lambda \mathbb{E}_{(x,y)\sim\mathcal{X}\times\mathcal{Y}} \left[\left(L\left(N_2(x), y\right) - L\left(N_3(x), x\right)\right)_+ \mathbb{1}_{\kappa_1(x)=N_2} + \left(L\left(N_1(x), y\right) - L\left(N_3(x), x\right)\right)_+ \mathbb{1}_{\kappa_1(x)=N_1} \right] \quad (7)$$

This problem can be reduced to a cost sensitive learning problem:

$$\min_{\kappa_1 \in \Gamma} \mathbb{E}_{(x,y)\sim\mathcal{X}\times\mathcal{Y}} \left[R_3(x,y)\mathbb{1}_{\kappa_1(x)=N_3} + R_2(x,y)\mathbb{1}_{\kappa_1(x)=N_2} + R_1(x,y)\mathbb{1}_{\kappa_1(x)=N_1} \right] + \tau(\kappa_1), \quad (8)$$

where the costs are defined:

$$R_1(x,y) = \left(L(N_1(x), y) - L(N_3(x), y)\right)_+$$
$$R_2(x,y) = \left(L(N_2(x), y) - L(N_3(x), y)\right)_+ + \tau(N_2)$$
$$R_3(x,y) = \tau(N_3).$$

5. Experimental Section

We evaluate our method on the Imagenet 2012 classification dataset (Russakovsky et al., 2015) which has 1000 object classes. We train using the 1.28 million training images and evaluate the system using 50k validation images. We use the pre-trained models from Caffe Model Zoo for Alexnet, GoogLeNet and Resnet50 (Krizhevsky et al., 2012; Szegedy et al., 2015; He et al., 2016). For preprocessing we follow the same routines proposed for these networks and verify the final network performances within a small margin ($< 0.1\%$).

Temporal measurements: We measure the network times using the built-in tool in the Caffe library on a server that utilizes a Nvidia Titan X Pascal with CuDNN 5. Since our focus is on the computational cost of the networks, we ignore the data loading and preprocessing times. The reported times are actual measurements including the policy overhead.

Policy meta-features: In addition to the outputs of the convolutional layers or earlier networks, we augment the feature space with the entropy of prediction probabilities. We experimented with pooled internal representations, however, empirical evidence shows that inclusion of the entropy feature with a simple linear policy outperforms significantly more complex policy functions excluding the entropy feature.

5.1. Network Selection

Baselines: Our full system, depicted in Figure 2, starts with Alexnet. Following the evaluation of Alexnet, the system determines for each example either to return the prediction, route the example to GoogLeNet, or route the example to Resnet50. For examples that are routed to GoogLeNet, the system either returns the prediction output by GoogLeNet or routes the example to Resnet50. As a baseline, we compare against a uniform policy and report performance from

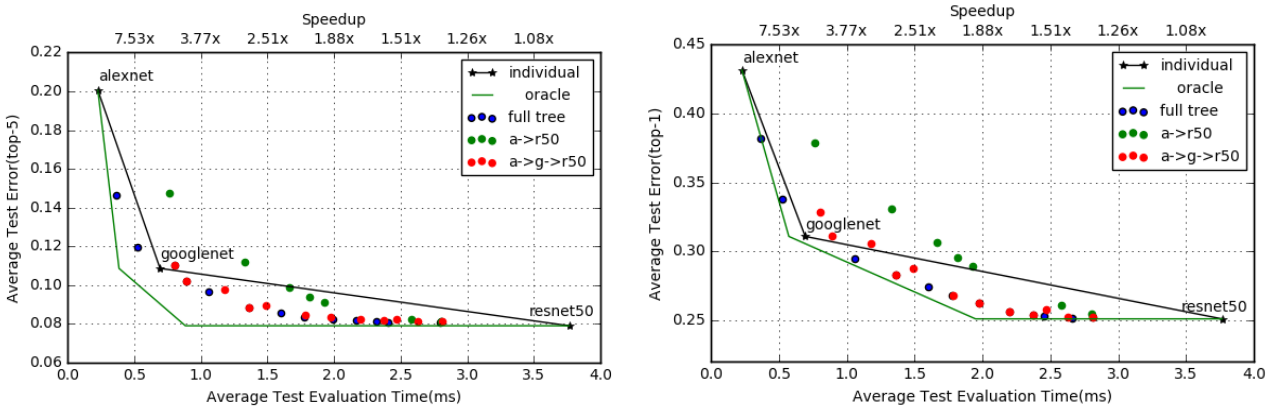

Figure 3. Performance of network selection policy on Imagenet (Left: top-5 error Right: top-1 error). Our full adaptive system (denoted with blue dots) significantly outperforms any individual network for almost all budget regions and is close to the performance of the oracle. The performances are reported on the validation set of ImageNet dataset.

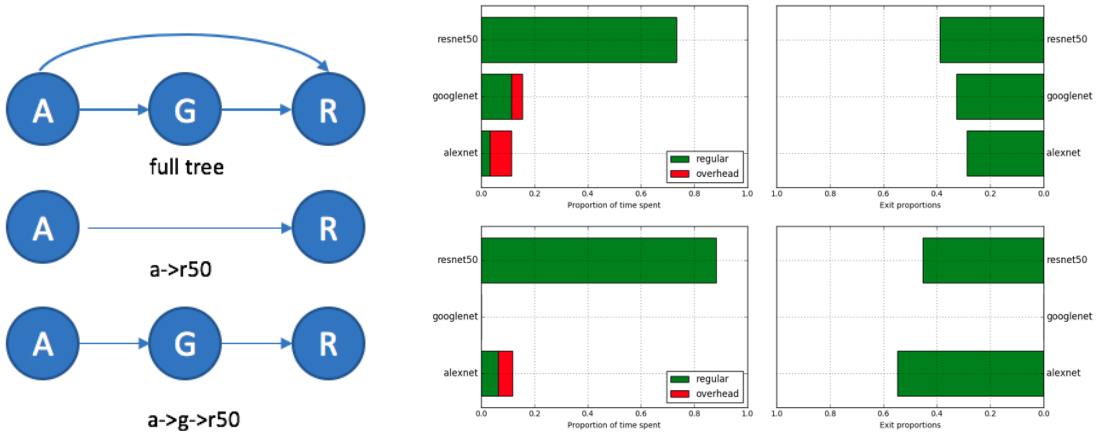

Figure 4. **(Left)** Different network selection topologies that we considered. Arrows denote possible jumps allowed to the policy. A, G and R denote Alexnet, GoogLeNet and Resnet50, respectively. **(Right)** Statistics for proportion of total time spent on different networks and proportion of samples that exit at each network. Top row is sampled at 2.0ms and bottom row is sampled at 2.8ms system evaluation

different system topologies. To provide a bound on the achievable performance, we show the peformance of a soft oracle. The soft oracle has access to classification labels and sends each example to the fastest model that correctly classifies the example. Since having access to the labels is too strong, we made the oracle softer by adding two constraints. First, it follows the same network topology, also it can not make decisions without observing the model feedback first, getting hit by the same overhead. Second, it can only exit at a cheaper model if all latter models agree on the true label. This second constraint is added due to the fact that our goal is not to improve the prediction performance of the system but to reduce the computational time, and therefore we prevent the oracle from "correcting" mistakes made by the most complex networks.

We sweep the cost trade-off parameter in the range 0.0 to 0.1 to achieve different budget points. Note that due to weights in our cost formulation, even when the pseudo labels are the same, policy behavior can greatly differ. Conceptually, the weights balance the importance of the samples that gain in classification loss in future stages versus samples that gain in computational savings by exiting early stages.

The results are demonstrated in Figure 3. We see that both full tree and a->g->r50 cascade achieve significant (2.8x) speedup over using Resnet50 while maintaining its accuracy within 1%. The classifier feedback for the policy has

Network	policy top-5	uniform top-5
GoogLeNet@1	9%	2%
GoogLeNet@2	22%	9%
GoogLeNet@5	33%	20%
Resnet50@1	8%	1%
Resnet50@2	18%	12%
Resnet50@5	22%	10%

Table 1. Early exit performances at different accuracy/budget trade-offs for different networks. @x denotes x loss from full model accuracy and reported numbers are percentage speed-ups.

a dramatic impact on its performance. Although, Alexnet introduces much less overhead compared to GoogLeNet (≈0.2 vs ≈0.7), the a->r50 policy performs significantly worse in lower budget regions. Our full tree policy learns to choose the best order for all budget regions. Furthermore, the policy matches the soft oracle performance in both the high and low budget regions.

Note that GoogLeNet is a very well positioned at 0.7ms per image budget, probably due to its efficiency oriented architectural design with inception blocks (Szegedy et al., 2015). For low budget regions, the overhead of the policy is a detriment, as even when it can learn to send almost half the samples to Alexnet instead of GoogLeNet with marginal loss in accuracy, the extra 0.23ms Alexnet overhead brings the balance point, $\approx 0.65ms$, very close to using only GoogLeNet at 0.7ms. The ratio between network evaluation times is a significant factor for our system. Fortunately, as mentioned before, for many applications the ratio of different models can be very high (cloud computing upload times, resnet versus Alexnet difference etc.).

5.2. Network Early Exits

Exit Architecture

To output a prediction following each convolutional layer, we train a single layer linear classifier after a global average pooling at every convolutional layer. We added global pooling to minimize the policy overhead in earlier exits.

For Resnet50 we added an exit after output layers of 2a, 2c, 3a, 3d, 4a and 4f. The dimensionality of the exit features after global average pooling are 256, 256, 512, 512, 1024 and 1024 in the same order as the layer names. For GoogLeNet we added the exits after concatenated outputs of every inception layer.

Table 1 shows the early exit performance for different networks. The gains are more marginal compared to network selection. We analyze this further in Section 5.4.

5.3. Network Selection Analysis

We further analyzed the network usage and runtime proportion statistics for samples at different budget regions. Fig. 4 demonstrates the results at three different budget levels. Full tree policy avoids using GoogLeNet altogether for high budget regions. This is the expected behavior since the a->r50 policy performs just as well in those regions and using GoogLeNet in the decision adds too much overhead. In mid level budget the policy distributes samples more evenly. Note that the sum of the overheads is close to useful runtime of cheaper networks in this region, but the system can afford that since the earlier networks are very lightweight.

5.4. Early Exits Analysis

In this section we analyze the performance of early exit layers. Interestingly, the accuracy gain per time is more linear within the same architecture compared to different network architectures. This explains why the adaptive policy works better for network selection compared to early exits.

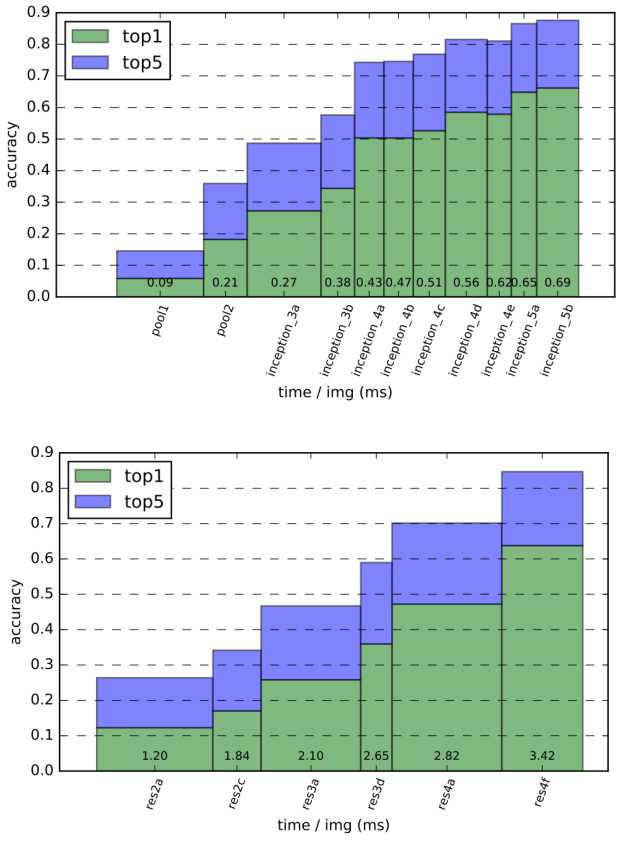

Figure 5. The plots show the accuracy gains at different layers for early exits for networks GoogLeNet (top) and Resnet50 (bottom).

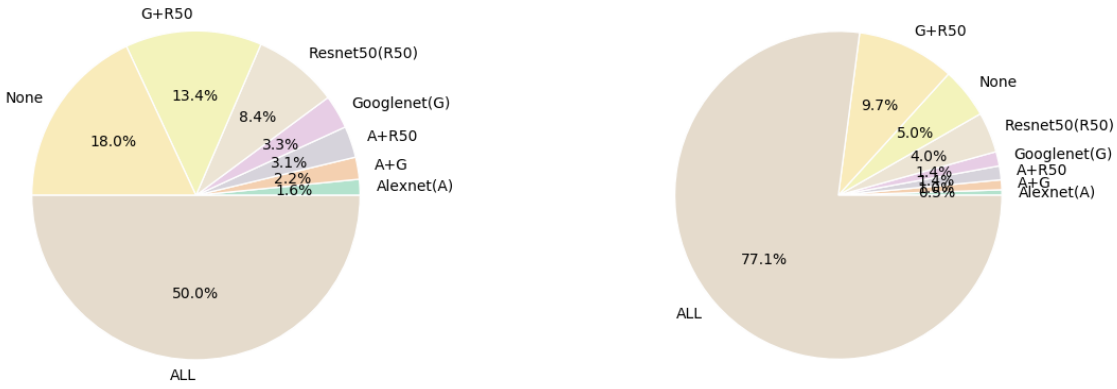

Figure 6. Analysis of top-1 and top-5 errors for different networks. Majority of the samples are easily classified by Alexnet, and only a minority of them require deeper networks.

5.5. Network Error Analysis

Figs. 6 and 6 show the distributions over examples of the networks that correctly label the example. Notably, 50% and 77% of the examples are correctly classified by all networks for top 1 and top 5 error, respectively. Similarly, 18% and 5% of the examples are incorrectly classified by all networks with respect to their top 1 and top 5 error, respectively. These results verify our hypothesis that for a large fraction of data, there is no need for costly networks. In particular, for the 68% and 82% of data with no change in top 1 and top 5 error, respectively, the use of any network apart from Alexnet is unnecessary and only adds unnecessary computational time.

Additionally, it is worth noting the balance between examples incorrectly classified by all networks, 18% and 5% respectively for top 1 and top 5 error, and the fraction of examples correctly classified by either GoogLeNet or Resnet but not Alexnet, 25.1% and 15.1% for top 1 and top 5 error, respectively. This behavior supports our observation that entropy of classification decisions is an important feature in making policy decisions, as examples likely to be incorrectly classified by Alexnet are likely to be classified correctly by a later network.

Note that our system is trained using the same data used to train the networks. Generally, the resulting evaluation error for each network on training data is significantly lower than error that arises on test data, and therefore our system is biased towards sending examples to more complex networks that generally show negligible training error. Practically, this problem is alleviated through the use of validation data to train the adaptive systems. In order to maintain the reported performance of the network without expansion of the training set, we instead utilize the same data for training both networks and adaptive systems, however we note that performance of our adaptive systems is generally better when trained on data excluded from the network training.

6. Conclusion

In this paper, we proposed two different schemes to adaptively trade off between model accuracy and model evaluation time for deep neural networks. We demonstrated that significant gains in computational time is possible through our novel policy with negligible loss in accuracy on ImageNet image recognition dataset. We believe that adaptivity is very important in the age of growing data and models with high variance in computational time and quality. We also showed that our method approximates an Oracle based policy that has benefit of access to true error for each instance from all the networks.

References

Bagherinezhad, Hessam, Rastegari, Mohammad, and Farhadi, Ali. Lcnn: Lookup-based convolutional neural network. *arXiv preprint arXiv:1611.06473*, 2016.

BuciluǍ, Cristian, Caruana, Rich, and Niculescu-Mizil, Alexandru. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 535–541. ACM, 2006.

Chen, Wenlin, Wilson, James T, Tyree, Stephen, Weinberger, Kilian Q, and Chen, Yixin. Compressing neural networks with the hashing trick. In *ICML*, pp. 2285–2294, 2015.

Courbariaux, Matthieu, Bengio, Yoshua, and David, Jean-Pierre. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in Neural Information Processing Systems*, pp. 3123–3131, 2015.

Figurnov, Mikhail, Ibraimova, Aizhan, Vetrov, Dmitry P, and Kohli, Pushmeet. Perforatedcnns: Acceleration

through elimination of redundant convolutions. In *Advances in Neural Information Processing Systems*, pp. 947–955, 2016.

Gong, Yunchao, Liu, Liu, Yang, Ming, and Bourdev, Lubomir. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*, 2014.

He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

Hinton, Geoffrey, Deng, Li, Yu, Dong, Dahl, George E, Mohamed, Abdel-rahman, Jaitly, Navdeep, Senior, Andrew, Vanhoucke, Vincent, Nguyen, Patrick, Sainath, Tara N, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29 (6):82–97, 2012.

Hinton, Geoffrey, Vinyals, Oriol, and Dean, Jeff. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

Hubara, Itay, Courbariaux, Matthieu, Soudry, Daniel, El-Yaniv, Ran, and Bengio, Yoshua. Binarized neural networks. In *Advances in Neural Information Processing Systems*, pp. 4107–4115, 2016a.

Hubara, Itay, Courbariaux, Matthieu, Soudry, Daniel, El-Yaniv, Ran, and Bengio, Yoshua. Quantized neural networks: Training neural networks with low precision weights and activations. *arXiv preprint arXiv:1609.07061*, 2016b.

Iandola, Forrest N, Han, Song, Moskewicz, Matthew W, Ashraf, Khalid, Dally, William J, and Keutzer, Kurt. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and< 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.

Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

Kusner, M, Chen, W, Zhou, Q, Xu, Z, Weinberger, K, and Chen, Y. Feature-cost sensitive learning with submodular trees of classifiers. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.

Liu, Baoyuan, Wang, Min, Foroosh, Hassan, Tappen, Marshall, and Pensky, Marianna. Sparse convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 806–814, 2015.

Rastegari, Mohammad, Ordonez, Vicente, Redmon, Joseph, and Farhadi, Ali. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, pp. 525–542. Springer, 2016.

Russakovsky, Olga, Deng, Jia, Su, Hao, Krause, Jonathan, Satheesh, Sanjeev, Ma, Sean, Huang, Zhiheng, Karpathy, Andrej, Khosla, Aditya, Bernstein, Michael, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

Sutskever, Ilya, Vinyals, Oriol, and Le, Quoc V. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112, 2014.

Szegedy, Christian, Liu, Wei, Jia, Yangqing, Sermanet, Pierre, Reed, Scott, Anguelov, Dragomir, Erhan, Dumitru, Vanhoucke, Vincent, and Rabinovich, Andrew. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–9, 2015.

Trapeznikov, K and Saligrama, V. Supervised sequential classification under budget constraints. In *International Conference on Artificial Intelligence and Statistics*, pp. 581–589, 2013.

Wang, J., Bolukbasi, T., Trapeznikov, K., and Saligrama, V. Model selection by linear programming. In *European Conference on Computer Vision*, pp. 647–662, 2014.

Wang, Joseph, Trapeznikov, Kirill, and Saligrama, Venkatesh. Efficient learning by directed acyclic graph for resource constrained prediction. In *Advances in Neural Information Processing Systems*, pp. 2152–2160, 2015.

Wen, Wei, Wu, Chunpeng, Wang, Yandan, Chen, Yiran, and Li, Hai. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 2074–2082, 2016.

Wu, Jiaxiang, Leng, Cong, Wang, Yuhang, Hu, Qinghao, and Cheng, Jian. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4820–4828, 2016.

Xu, Z., Kusner, M., Chen, M., and Weinberger, K. Cost-sensitive tree of classifiers. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 133–141, 2013.